

Remarks

I. Priority

The Office Action asserts that applicants have not complied with one or more conditions for receiving the benefit of an earlier filing date under 35 U.S.C. 119 or 120, stating:

The later-filed application must be an application for a patent for an invention which is also disclosed in the prior application (the parent or original nonprovisional application or provisional application). The disclosure of the invention in the parent application and in the later-filed application must be sufficient to comply with the requirements of the first paragraph of 35 U.S.C. 112. See *Transco Products, Inc. v. Performance Contracting, Inc.*, 38 F.3d 551, 32 USPQ2d 1077 (Fed. Cir. 1994).

The disclosure of the prior-filed application, Application No. 10/005,536 (or any of the other prior applications for which benefit is claimed under 35 U.S.C. 119 or 120) fails to provide adequate support or enablement in the manner provided by the first paragraph of 35 U.S.C. 112 for one or more claims of this application.

Applicant's prior applications do not provide adequate disclosure for the subject matter pertaining to the "hybrid storage area" or the "re-assembly storage area" as claimed (i.e. all limitations regarding these storage areas as well as the claimed functionality).

MPEP 201.07 recites, "The disclosure presented in the continuation must be the same as that of the original application." It appears that the disclosure for this specification contains subject matter not contained in the prior applications.

Accordingly, claims 1-40 are not entitled to the benefit of the prior applications.

Applicants respectfully disagree with the Examiner's denial of the benefit of prior applications for claims 1-40. Initially, applicants note that the case relied upon by the Examiner, *Transco*, involved the issue of whether the "best mode" needs to be updated upon the filing of a "continuing application," not whether an application was entitled to the benefit of the filing date of a prior-filed application. As to the latter issue, *dicta* in *Transco* merely requires "common subject matter" for such entitlement. In general, "the test for sufficiency of support in a parent application is whether the disclosure of the application relied upon 'reasonably conveys to the artisan that the inventor had possession at that time of the later claimed subject matter.' *In re Kaslow*, 707 F.2d 1366, 1375, 217 U.S.P.Q. (BNA) 1089, 1096 (Fed. Cir. 1983)." *Ralston Purina Co. v. Far-Mar-Co.*, 772 F.2d 1570, 1575 (Fed. Cir. 1985). Moreover, cases such as *KangarooOS*,

U.S.A., Inc. v. Caldor, Inc., 778 F.2d 1571, 1574 (Fed. Cir. 1985) hold that a utility patent may claim priority from a design patent, even though no claims existed in the design patent other than claiming the drawings that were shown. Similarly, cases such as *In re Hogan*, 559 F.2d 595, 608 (C.C.P.A. 1977) hold that “claimed subject matter need not be described *in haec verba* in the application to satisfy the written-description-of-the-invention requirement.”

Indeed, *Hogan* involved a situation where a broad claim ostensibly covered later enabled species even though those species were not enabled in the original application. Although *Hogan* involved chemical species that are much less predictable than mechanical or electrical inventions, *Hogan* held that failure to enable those species did not undermine the claim to the filing date the earlier application. *Hogan* at 606. “To now say that appellants should have disclosed in 1953 the amorphous form which on this record did not exist until 1962, would be to impose an impossible burden on inventors and thus on the patent system. There cannot, in an effective patent system, be such a burden placed on the right to broad claims.” *Id.*

Applicants therefore respectfully disagree with the Office Action assertion that “prior applications do not provide adequate disclosure for the subject matter pertaining to the ‘hybrid storage area’ or the ‘re-assembly storage area’ as claimed (i.e. all limitations regarding these storage areas as well as the claimed functionality).” While it is true that applicants did not use the exact *nonce words* “hybrid storage area” or “re-assembly storage area” that were coined in U.S. Patent No. 6,480,489 to Muller et al. (“hereinafter “Muller”), it is clear that applicants’ parent cases disclosed the subject matter that those words represent.

Perhaps because the nonce words coined by Muller are not terms of art, column 58, lines 8-11 of that reference explains: “A buffer used to store portions of more than one type of packet-such as a header buffer used to store headers and small packets, or a non-re-assembly buffer used to store MTU and jumbo packets-may be termed a ‘hybrid’ buffer.” Similarly, column 4, lines 48-50 of that reference states: “A re-assembly buffer may be used to re-assemble data from multiple packets of a single communication flow.” In concert with this explanation of those nonce words is applicants’ claim 37, which recites in part: “a re-assembly storage area configured to store data portions of a plurality

of packets from a single communication flow;” and “a hybrid storage area configured to store: header portions of the plurality of packets; and one or more packets smaller than a first predetermined size.”

Ample support for these limitations can be found in applicants’ early disclosures. For example, regarding the limitation of “a hybrid storage area,” that is “configured to store: header portions of the plurality of packets; and one or more packets smaller than a first predetermined size,” applicants reproduce below pertinent portions of pages 9 – 13 of application Serial No. 60/098,296 (hereinafter “the ‘296 app.”), with added emphasis and parenthetical insertion of explanatory or nonce words. Support for the limitation of a “re-assembly storage area” that is “configured to store data portions of a plurality of packets from a single communication flow;” can also be found in these portions of the ‘296 app.

2.3.2. Support small and large buffers on receive

In order to reduce further the number of writes to the INIC, and to reduce the amount of memory being used by the host, we support two different buffer sizes. A small (*hybrid*) buffer contains roughly 200 bytes of data payload, as well as extra fields containing status about the received data bringing the total size to 256 bytes. We can therefore pass 16 of these small buffers at a time to the INIC. Large (*data or re-assembly*) buffers are 2k in size. *They are used to contain any fast or slow-path data that does not fit in a small buffer.* Note that when we have a large fast-path receive, a small (*hybrid*) buffer will be used to indicate a small piece (*header portion*) of the data, while *the remainder of the data* will be DMA’d directly into *memory (data buffers or re-assembly storage area)*.

...

2.4.1. Fast-path 56k NetBIOS session message

Let’s say a 56k NetBIOS session message is received on the INIC. The first segment will contain the NetBIOS header, which contains the total NetBIOS length. A small chunk (*header portion*) of this first segment is provided to the host by filling in a small (*hybrid*) receive buffer, modifying the interrupt status register on the host, and raising the appropriate interrupt line. Upon receiving the interrupt, the host will read the ISR, clear it by writing back to the INIC’s Interrupt Clear Register, and will then process its small receive buffer queue looking for receive buffers to be processed. Upon finding the small buffer, it will indicate the small amount of data up to the client to be processed by NetBIOS. It will also, if necessary, replenish the receive buffer pool on the INIC by passing off a pages worth of small buffers. Meanwhile, the NetBIOS client will allocate a memory pool (*data buffers or re-assembly storage area*) large enough to hold the entire NetBIOS message, and will pass this address or

set of addresses down to the transport driver. The transport driver will allocate an INIC command buffer, fill it in with the list of addresses, set the command type to tell the INIC that this is where to put the receive data, and then pass the command off to the INIC by writing to the command register. When the INIC receives the command buffer, it will DMA the remainder of the NetBIOS data, as it is received, into the memory address or addresses (*data buffers or re-assembly storage area*) designated by the host. Once the entire NetBIOS transaction is complete, the INIC will complete the command by writing to the response buffer with the appropriate status and command buffer identifier.

...

2.4.2. Slow-path receive

If the INIC receives a frame that does not contain a TCP segment for one of its CCB's, it simply passes it to the host as if it were a dumb NIC. *If the frame fits into a small (hybrid) buffer (~200 bytes or less), then it simply fills in the small buffer with the data and notifies the host.* Otherwise it places the data in a large buffer, writes the address of the large buffer into a small buffer, and again notifies the host. The host, having received the interrupt and found the completed small buffer, checks to see if the data is contained in the small buffer, and if not, locates the large buffer. Having found the data, the host will then pass the frame upstream to be processed by the standard protocol stack. It must also replenish the INIC's small and large receive buffer pool if necessary.

...

We take advantage of this fact by allocating large and small receive buffers. *If a received frame fits in a small buffer, the INIC will use a small (hybrid) buffer. Otherwise it will use a large (data) buffer.* A problem with that system then is preserving receive order. If we were to maintain a small and a large buffer queue, there would be no way to know in which order two frames, one small and one large, were received. A solution is to maintain a single receive queue of small buffers. We pass the small buffers in blocks of 16 at a time to the INIC, and they are guaranteed to be returned to us in the order in which they were given to the INIC. The small buffer contains status about the receive as well as small frames. If a received frame does not fit in the small buffer, then we allocate a large buffer and place a pointer to that large buffer in the small buffer. Thus, large buffers are only returned to the driver when attached to small buffers.

The remainder of this section covers this in greater detail.

3.2. Receive Interface

3.2.1. Receive Interface Overview

As mentioned above, the fast-path flow puts a *header* into a header (*hybrid*) buffer that is then forwarded to the host. The host uses the header to determine what further data is following, allocates the necessary host buffers (*data buffers or re-assembly storage area*), and these are passed back to the INIC via a command to the INIC. The INIC then fills these

buffers from data it was accumulating on the card and notifies the host by sending a response to the command. *Alternatively, the fast-path may receive a header and data that is a complete request, but that is also too large for a header (hybrid) buffer. This results in a header and data buffer being passed to the host.* This latter flow is identical to the slow-path flow which also puts all the data into the header (*hybrid*) buffer or, if the header buffer is too small, uses a large (2K) host (*data*) buffer for all the data.

Although such support was detailed in the Request to Provoke Interference, in the chart below applicants provide examples of support for each claim limitation, along with a parenthetical listing of the page and line numbers where such support can be found in the '296 app.

Claim Element	Support in the '296 app. (Page #: Line #)
1. A method of storing a portion of a packet in a host computer memory, comprising:	
receiving a first packet at a communication interface;	(6:30-31) "When a frame is received by the INIC , it must verify it completely before it even determines whether it belongs to one of its CCBs or not."
receiving a second packet at said communication interface;	(38:17) "4.7.2.2. Two types of receive packets "
storing a header portion of said first packet in a hybrid storage area of a host computer;	(13:17-18) "As mentioned above, the fast-path flow puts a header into a header buffer that is then forwarded to the host." (13:4-5) "If a received frame fits in a small buffer , the INIC will use a small buffer. Otherwise it will use a large buffer."
if said first packet includes a data portion, storing said data portion in a data storage area of said host computer; and	(5:27-28) "In the fast path case, network data is given to the host after the headers have been processed and stripped." (13:17-21) "As mentioned above, the fast-path flow puts a header into a header buffer that is then forwarded to the host. The host uses the header to determine what further data is following, allocates the necessary host buffers, and these are passed back to the INIC via a command to the INIC. The INIC then fills these buffers from data it was accumulating on the card..."
if said second packet is smaller than a predetermined size, storing said second packet in said hybrid storage area.	(13:4-5) "If a received frame fits in a small buffer , the INIC will use a small buffer. Otherwise it will use a large buffer."

<p>2. The method of claim 1, wherein said hybrid storage area is substantially equal in size to one memory page of said host computer.</p>	<p>(10:3-4) <i>"Like the small and large receive buffers, a page worth of response buffers is passed to the INIC at a time."</i> (10:3-4) "It will also, if necessary, replenish the receive buffer pool on the INIC by passing off <i>a pages worth of small buffers.</i>" (12:42-44) "In the driver we allocate <i>a block of contiguous memory (typically a page, which is typically 4k).</i>"</p>
<p>3. The method of claim 1, wherein said predetermined size is approximately 256 bytes.</p>	<p>(14:6-7) <i>"Header buffers in host memory are 256 bytes long, and are aligned on 256 byte boundaries."</i></p>
<p>4. The method of claim 1, further comprising:</p> <p>receiving a code associated with said first packet, said code indicating that a data portion of said first packet is re-assembleable;</p> <p>wherein said data storage area is a re-assembly data storage area.</p>	<p>(6:36-38) "The <i>header is fully parsed</i> by hardware and its <i>type is summarized</i> in a single status word. (41:25-27) "2. If the <i>type field contains our custom INIC type</i> (TCP for example): A. If the <i>header buffer specifies a fast-path connection</i>, allocate one or more mbufs headers to map the header and possibly data buffers." (13:17-21) "As mentioned above, the fast-path flow puts a header into a header buffer that is then forwarded to the host. The host uses the header to determine what further data is following, <i>allocates the necessary host buffers</i>, and these are passed back to the INIC via a command to the INIC. The <i>INIC then fills these buffers from data</i> it was accumulating on the card and notifies the host by sending a response to the command."</p>
<p>5. The method of claim 1, further comprising:</p> <p>receiving a code associated with said second packet, said code indicating that a data portion of said second packet is not re-assembleable;</p> <p>wherein said data storage area is a non-re-assembly data storage area.</p>	<p>(85:26-30) "As frames are received, they are placed into 2K-byte DRAM buffers by the Receive hardware sequencer, along with 16 bytes of the above frame status. A pointer to the last byte + 1 of this buffer is queued into the Q_RECV queue. The <i>pointer contains a bit (bit 29) that informs the microcode if this frame is definitely not a fast-path candidate...</i>" (85:38) "If bit 29 is set, this frame is going <i>slow-path.</i>" (87:9-15) "The following is a summary of <i>slow-path processing</i>:" <ul style="list-style-type: none"> • Examine frame status bytes to determine if frame is in-error; if so, only these status bytes will be sent to the host. • <i>Move the frame into either a small or a large host buffer</i> via DMA. It is not split across these buffers. • Set frame status and address details and DMA status to the host. • Send event to the Utility processor to post Receive status in the ISR." </p>

<p>6. The method of claim 1, further comprising receiving a code with said first packet, said code indicating that said first packet does not contain a data portion.</p>	<p>(14:6-7) "There will be a field in the header buffer indicating it has valid data." (86:23-29) "Examine the frame header to generate an event from it. The Receive events that can be generated on a given context from a frame are: <i>receive a pure ACK...</i> <i>receive a window probe...</i>"</p>
<p>7. The method of claim 1, further comprising receiving a code with said second packet, said code indicating that said second packet is smaller than said predetermined size.</p>	<p>(70:12-13) "Receive data is passed from the INIC to the host by filling in a header buffer. The header buffer contains information about the data, such as the length."</p>
<p>8. The method of claim 1, further comprising padding said hybrid storage area to align said header portion with a predetermined index in said hybrid storage area.</p>	<p>(14:6-7) "Header buffers in host memory are 256 bytes long, and are aligned on 256 byte boundaries."</p>
<p>9. The method of claim 1, wherein said first packet belongs to a first communication flow and said second packet belongs to a second communication flow.</p>	<p>(6:10-16) "This introduces the notion of a Communication Control Block (CCB) cache. A CCB is a structure that contains the entire context associated with a connection. This includes the source and destination IP addresses and source and destination TCP ports that define the connection. It also contains information about the connection itself such as the current send and receive sequence numbers, and the first-hop MAC address, etc. The complete set of CCBs exists in host memory, but a subset of these may be "owned" by the card at any given time. This subset is the CCB cache. The INIC can own up to 256 CCBs at any given time."</p>
<p>10. A method of transferring multiple packets in a communication flow into a host computer, comprising:</p> <p>receiving a first packet at a network interface for transfer to a host computer;</p> <p>identifying a communication flow comprising said first packet;</p>	<p>(6:30-31) "When a frame is received by the INIC, it must verify it completely before it even determines whether it belongs to one of its CCBs or not."</p> <p>(86:1-8) "The receive sequencer has already generated a hash based on the network and transport addresses, e.g., IP source and destination addresses and TCP ports. This hash is used to index directly into a hash table on the INIC that points to entries in a CCB header table. The header table entries are chained on the hash table entry. The microcode uses the hash to determine if a CCB exists on the INIC for this frame. It does this by following this chain from the hash</p>

<p>receiving a second packet at said network interface, wherein said communication flow further comprises said second packet;</p> <p>associating a code with said first packet to indicate whether a data portion of said first packet is re-assembleable with a data portion of another packet in said communication flow;</p> <p>receiving a set of descriptors from the host computer, wherein an aggregate size of said set of descriptors approximates a cache line size of the host computer;</p>	<p>table entry, and for each chained header table entry, comparing its source and destination addresses and ports with those of the frame."</p> <p>(6:10-12) "A CCB is a structure that contains the entire context associated with a connection. This includes the source and destination IP addresses and source and destination TCP ports that define the connection.</p> <p>(4:16-19) "The ... context is ... identified by the IP source and destination addresses and TCP source and destination ports."</p> <p>See above.</p> <p>(85:26-30) "As frames are received, they are placed into 2K-byte DRAM buffers by the Receive hardware sequencer, along with 16 bytes of the above frame status. A pointer to the last byte + 1 of this buffer is queued into the Q_RECV queue. The pointer contains a bit (bit 29) that informs the microcode if this frame is definitely not a fast-path candidate (e.g., not TCPIP, or has an error of some sort)."</p> <p>(86:1-9) "If bit 29 is not set, then there may be an onboard CCB for this frame. The receive sequencer has already generated a hash based on the network and transport addresses, e.g., IP source and destination addresses and TCP ports. This hash is used to index directly into a hash table on the INIC that points to entries in a CCB header table. The header table entries are chained on the hash table entry. The microcode uses the hash to determine if a CCB exists on the INIC for this frame. It does this by following this chain from the hash table entry, and for each chained header table entry, comparing its source and destination addresses and ports with those of the frame. If a match is found, then the frame will be processed against the CCB by the INIC. If not, then the frame is sent for slow-path processing."</p> <p>(145:5-1) "The read multiple command requires that the Pmi sequencer be capable of transferring a cache line or more of data. To accomplish this end, Pmi will automatically perform partial cache line bursts until it has aligned the transfers on a cache line boundary at which time it will begin usage of the read multiple command. The Sram fifo depth, of 256 bytes, has been chosen in order to allow Pmi to accommodate cache line sizes up to 128 bytes. Provided the cache line size is less than 128 bytes, Pmi will perform multiple, contiguous cache line bursts until it has filled the fifo."</p> <p>(12:41-45) "Thus we needed an efficient way in which to pass receive buffer addresses to the INIC. We accomplished this by passing a block of receive buffers to the INIC at one time. In the driver we allocate a</p>
---	--

<p>storing a header portion of said first packet in a header storage area;</p> <p>storing said data portion of said first packet in a re-assembly storage area; and</p> <p>storing a data portion of said second packet in said re-assembly storage area.</p>	<p>block of contiguous memory (typically a page, which is typically 4k). We write the address of that block to the INIC with the bottom bits of the address specifying the number of buffers in the block." (14:6-13) "Header buffers in host memory are 256 bytes long, and are aligned on 256 byte boundaries. There will be a field in the header buffer indicating it has valid data. This field will initially be reset by the host before passing the buffer descriptor to the INIC. A set of header buffers are passed from the host to the INIC by the host writing to the Header Buffer Address Register on the INIC. This register is defined as follows:</p> <p>Bits 31-8 Physical address in host memory of the first of a set of contiguous header buffers</p> <p>Bits 7-0 Number of header buffers passed."</p> <p>(13:17-18) "As mentioned above, the fast-path flow puts a header into a header buffer that is then forwarded to the host."</p> <p>(13:18-21) "The host uses the header to determine what further data is following, allocates the necessary host buffers, and these are passed back to the INIC via a command to the INIC. The INIC then fills these buffers from data it was accumulating on the card and notifies the host by sending a response to the command."</p> <p>(10:25-32) "the NetBIOS client will allocate a memory pool large enough to hold the entire NetBIOS message, and will pass this address or set of addresses down to the transport driver. The transport driver will allocate an INIC command buffer, fill it in with the list of addresses, set the command type to tell the INIC that this is where to put the receive data, and then pass the command off to the INIC by writing to the command register. When the INIC receives the command buffer, it will DMA the remainder of the NetBIOS data, as it is received, into the memory address or addresses designated by the host."</p> <p>See above.</p>
<p>11. The method of claim 10, wherein said re-assembly storage area is substantially equal in size to one memory page of said host computer.</p>	<p>(9:23-24) "In order to avoid having to write to the INIC for every receive frame, we instead allow the host to pass off a pages worth (4k) of buffers in a single write."</p> <p>(9:30-37) "Large buffers are 2k in size. They are used to contain any fast or slow-path data that does not fit in a small buffer. Note that when we have a large fast-path receive, a small buffer will be used to indicate a small piece of the data, while the remainder of the data will be DMA'd directly into memory. Large buffers are never passed to the host by themselves, instead they are always accompanied by a small buffer which contains status about the receive along with the large buffer</p>

	address. By operating in the manner, the driver must only maintain and process the small buffer queue. Large buffers are returned to the host by virtue of being attached to small buffers. Since large buffers are 2k in size they are passed to the INIC 2 buffers at a time."
12. The method of claim 10, wherein said re-assembly storage area is used to store only data portions of packets in said communication flow.	(5:27-28) "In the fast path case, network data is given to the host after the headers have been processed and stripped. "
13. The method of claim 10, further comprising retrieving an identifier of a first storage area in a host computer memory.	(7:23-27) "We will make use of this feature by providing a small amount of any received data to the host, with a notification that we have more data pending. When this small amount of data is passed up to the client, and it returns with the address in which to put the remainder of the data , our host transport driver will pass that address to the INIC which will DMA the remainder of the data into its final destination." (21:33-41) "In the "large data input" case, where "bytes available" exceeds the packet length, the TDI client then provides an MDL, associated with an IRP, which must be completed when this MDL is filled. (This IRP/MDL may come back either in the response to ATCP's call of the receive handler, or as an explicit TDI_RECEIVE request.) The ATCP driver builds a "receive request" from the MDL information, and passes this to the INIC. This request contains: The TCP context identifier. Size and offset information. A scatter/gather list of physical addresses corresponding to the MDL pages. "
14. The method of claim 13, further comprising: placing said first storage area identifier in a data structure configured to hold storage area identifiers; wherein said first storage area identifier is identifiable by an index in said data structure.	(14:27-36) "Receive data buffers are allocated in blocks of 2, 2k bytes each (4k page). In order to pass receive data buffers to the INIC, the host must write two values to the INIC. The first value to be written is the Data Buffer Handle . The buffer handle is not significant to the INIC, but will be copied back to the host to return the buffer to the host. The second value written is the Data Buffer Address . This is the physical address of the data buffer. When both values have been written, the INIC will add these values to FreeType queue of data buffer descriptors . The INIC will extract 2 entries each time when dequeuing. Data buffers will be allocated and used by the INIC as needed. For each data buffer used, the data buffer handle will be copied into a header buffer. Then the header buffer will be returned to the host."
15. The method of claim 14, further comprising using said index to identify said first storage area for storing a portion	(14:34-36) "Data buffers will be allocated and used by the INIC as needed. For each data buffer used, the data buffer handle will be copied into a header buffer . Then the header buffer will be returned to the host."

of said first packet.	
16. The method of claim 14, wherein said first storage area comprises one of said header storage area and said re-assembly storage area, the method further comprising: configuring a descriptor to store said index of said first storage area identifier to inform said host computer of the use of said first storage area to store one of said header portion and said data portion.	(14:15-17) "For each interface, the INIC will maintain a queue of these header descriptors in the SmallHType queue in it's own local memory, adding to the end of the queue every time the host writes to one of the Header Buffer Address Registers."
18. The method of claim 9, further comprising parsing a header portion of said first packet.	(6:36-37) "The header is fully parsed by hardware and its type is summarized in a single status word."
19. The method of claim 18, wherein said parsing comprises: retrieving an identifier of a source of said first packet; and retrieving an identifier of a destination of said first packet.	(6:36-38) "The header is fully parsed by hardware and its type is summarized in a single status word. The checksum is also verified automatically in hardware, and a hash key is created out of the IP addresses and TCP ports to expedite CCB lookup." (160:6-7) "The Mac, network, transport and session information is analyzed as each byte is received and stored in the assembly register (AssyReg)." (86:1-3) "The receive sequencer has already generated a hash based on the network and transport addresses, e.g., IP source and destination addresses and TCP ports ."
20. The method of claim 19, wherein said identifying a communication flow comprises assembling said source identifier and said destination identifier to form a flow key.	(6:10-14) " CCBs are initialized by the host during TCP connection setup ." (6:10-12) "A CCB is a structure that contains the entire context associated with a connection . This includes the source and destination IP addresses and source and destination TCP ports that define the connection ." (4:16-19) "The ... context is ... identified by the IP source and destination addresses and TCP source and destination ports ." (86:1-8) "The receive sequencer has already generated a hash based on the network and transport addresses, e.g., IP source and destination addresses and TCP ports . This hash is used to index directly into a hash table on the INIC that points to entries in a CCB header table. The header table entries are chained on the hash table entry. The microcode uses the hash to determine if a CCB exists on the

	INIC for this frame. It does this by following this chain from the hash table entry, and for each chained header table entry, comparing its source and destination addresses and ports with those of the frame."
21. The method of claim 20, said identifying further comprising obtaining an index of said flow key in a database of flow keys.	(86:1-8) "The receive sequencer has already generated a hash based on the network and transport addresses, e.g., IP source and destination addresses and TCP ports . This hash is used to index directly into a hash table on the INIC that points to entries in a CCB header table . The header table entries are chained on the hash table entry. The microcode uses the hash to determine if a CCB exists on the INIC for this frame. It does this by following this chain from the hash table entry, and for each chained header table entry, comparing its source and destination addresses and ports with those of the frame." (24:36-39) "The initial command from ATCP to INIC expresses an "intention" to hand out the context. It carries a context number , context numbers are allocated by the ATCP driver, which keeps a per-INIC table of free and in-use context numbers . It also includes the source and destination IP addresses and ports , which will allow the INIC to establish a "provisional" context."
22. The method of claim 9, wherein said identifying comprises: receiving a virtual connection identifier, wherein a communication flow comprising said first packet can be identified by said virtual connection identifier.	 24:36-39) "The initial command from ATCP to INIC expresses an "intention" to hand out the context. It carries a context number , context numbers are allocated by the ATCP driver, which keeps a per-INIC table of free and in-use context numbers . It also includes the source and destination IP addresses and ports ..."
23. The method of claim 9, wherein said storing a header portion comprises: retrieving an index of a header storage area of a host computer; retrieving an identifier of a location in said header storage area; and storing a header portion of said first packet at said location in said header storage area.	 (14:9-11) "A set of header buffers are passed from the host to the INIC by the host writing to the Header Buffer Address Register on the INIC. (14:11-14) "This register is defined as follows: Bits 31-8 Physical address in host memory of the first of a set of contiguous header buffers Bits 7-0 Number of header buffers passed." See above.
24. The method of claim 23, wherein said header storage area index comprises an index	(14:28-30) "In order to pass receive data buffers to the INIC, the host must write two values to the INIC. The first value to be written is the Data Buffer Handle. The buffer handle is not significant to the INIC, but

of said header storage area within a collection of storage area identifiers.	will be copied back to the host to return the buffer to the host. The second value written is the Data Buffer Address. This is the physical address of the data buffer. When both values have been written, the INIC will add these values to FreeType queue of data buffer descriptors. The INIC will extract 2 entries each time when dequeuing."
25. The method of claim 24, wherein said identifier of a location in said header storage area comprises an address within said header storage area.	(14:12-14) "Bits 31-8 Physical address in host memory of the first of a set of contiguous header buffers."
26. The method of claim 9, wherein said storing said data portion comprises:	
retrieving an identifier of a re-assembly storage area of a host computer;	(7:24-26) "When this small amount of data is passed up to the client, and it returns with the address in which to put the remainder of the data , our host transport driver will pass that address to the INIC..."
retrieving an identifier of a location in said re-assembly storage area; and	(10:25-32) "Meanwhile, the NetBIOS client will allocate a memory pool large enough to hold the entire NetBIOS message, and will pass this address or set of addresses down to the transport driver. The transport driver will allocate an INIC command buffer, fill it in with the list of addresses , set the command type to tell the INIC that this is where to put the receive data, and then pass the command off to the INIC by writing to the command register.
storing said data portion of said first packet at said location in said re-assembly storage area.	(7:24-27) "When this small amount of data is passed up to the client, and it returns with the address in which to put the remainder of the data, our host transport driver will pass that address to the INIC which will DMA the remainder of the data into its final destination ."
27. The method of claim 26, wherein said retrieving an identifier of a re-assembly storage area comprises reading an entry in a re-assembly storage table, said entry corresponding to said communication flow.	(21:37-39) "The ATCP driver builds a "receive request" from the MDL information, and passes this to the INIC. This request contains: The TCP context identifier . Size and offset information. A scatter/gather list of physical addresses corresponding to the MDL pages."
28. The method of claim 18, wherein said data portion is of unknown size prior to said parsing.	(162:6-10) "27:00 address Represents the last address +1 to which frame data was transferred. This can be used to determine the size of the frame received by first subtracting one from the address then comparing with the buffer boundary as indicated by the size bits." (165:20) "14:00 LengthCnt Refer to E110 Technical Manual."
29. A network interface for transferring a packet from a network to a host computer,	

<p>comprising:</p> <p>a parser configured to examine one or more packets received from a network;</p> <p>a flow manager configured to manage a first communication flow comprising said one or more packets;</p> <p>a header storage area configured to store header portions of said one or more packets;</p> <p>a first re-assembly storage area configured to store data portions of said one or more packets;</p> <p>a second re-assembly storage area; and</p> <p>a transfer module configured to:</p> <p>transfer a data portion of a first packet of a second communication flow into said second re-assembly</p>	<p>(6:36-38) "The header is fully parsed by hardware and its type is summarized in a single status word. The checksum is also verified automatically in hardware, and a hash key is created out of the IP addresses and TCP ports to expedite CCB lookup."</p> <p>(8:27-41) "As frames are received, they are placed into 2K-byte DRAM buffers by the Receive hardware sequencer, along with 16 bytes of the above frame status. A pointer to the last byte + 1 of this buffer is queued into the Q_RECV queue. The pointer contains a bit (bit 29) that informs the microcode if this frame is definitely not a fast-path candidate (e.g., not TCPIP, or has an error of some sort). Receive frame processing involves extracting this pointer from the Receive hardware queue, and setting up a DMA into an SRAM header buffer of the first X bytes from the DRAM frame buffer. The size of the DMA is determined by whether bit 29 is set or not. If it is set (this frame is not a fast-path candidate), then only the status bytes are needed by the microcode, so the size would be 16 bytes. Otherwise up to 92 bytes are DMA'd – sufficient to get all useful headers. When this DMA is complete, the status bytes are used by the microcode to determine whether to jump to fast-path or slow-path processing."</p> <p>(13:17-18) "As mentioned above, the fast-path flow puts a header into a header buffer that is then forwarded to the host."</p> <p>(5:27-28) "In the fast path case, network data is given to the host after the headers have been processed and stripped."</p> <p>(13:17-21) "As mentioned above, the fast-path flow puts a header into a header buffer that is then forwarded to the host. The host uses the header to determine what further data is following, allocates the necessary host buffers, and these are passed back to the INIC via a command to the INIC. The INIC then fills these buffers from data it was accumulating on the card..."</p> <p>See above.</p> <p>(13:20-21) "Alternatively, the fast-path may receive a header and data that is a complete request, but that is also too large for a header buffer. This results in a header and data buffer being passed to the host."</p> <p>(14:1-3) "Note that the order in which data is written is important. Data</p>
--	--

<p>storage area if said first packet is larger than a predetermined size; and</p> <p>otherwise, transfer said first packet into said header storage area.</p>	<p>buffers are moved by DMA into the host before the header buffer, since the header buffer contains the status word designating that the data has arrived."</p> <p>(7:28-38) "Clearly there are circumstances in which this does not make sense. When a small amount of data (500 bytes for example), with a push flag set indicating that the data must be delivered to the client immediately, it does not make sense to deliver some of the data directly while waiting for the list of addresses to DMA the rest. Under these circumstances, it makes more sense to deliver the 500 bytes directly to the host, and allow the host to copy it into its final destination. While various ranges are feasible, it is currently preferred that anything less than a segment's (1500 bytes) worth of data will be delivered directly to the host, while anything more will be delivered as a small piece (which may be 128 bytes), while waiting until receiving the destination memory address before moving the rest."</p>
<p>30. A computer readable storage medium storing instructions that, when executed by a computer, cause the computer to perform a method of transferring a packet from a network to a host computer, the method comprising:</p> <p>receiving a first packet at a communication interface;</p> <p>receiving a second packet at said communication interface;</p> <p>storing a header portion of said first packet in a hybrid storage area of a host computer;</p>	<p>(110:12) "The processor instructions reside in the on-chip control-store, which is implemented as a mixture of ROM and Sram."</p> <p>(80:6-31) "As specified in other sections, the INIC supplies a set of 3 custom processors (CPUs) that provide considerable hardware-assist to the microcode running thereon...</p> <ul style="list-style-type: none"> • Multiple register contexts or process slots with register access controlled by simply setting a process register. The Protocol Processor will provide 512 SRAM-based registers to be shared among the 3 CPUs in any way desired. The current implementation uses 16 processes of 16 registers each, leaving 256 scratch registers to be shared. • A set of CPU-specific registers that are the same local-cpu register number, but for which the real register is determined by an offset based on the CPU number; this allows multiple CPUs to execute the same code at the same time without register clashes or interlocks. These registers are a part of the above-mentioned scratch pool." <p>(6:30-31) "When a frame is received by the INIC, it must verify it completely before it even determines whether it belongs to one of its CCBs or not."</p> <p>(38:17) "4.7.2.2. Two types of receive packets"</p> <p>(13:17-18) "As mentioned above, the fast-path flow puts a header into a header buffer that is then forwarded to the host."</p> <p>(13:4-5) "If a received frame fits in a small buffer, the INIC will use a small buffer. Otherwise it will use a large buffer."</p>

<p>if said first packet includes a data portion, storing said data portion in a data storage area of said host computer; and</p> <p>if said second packet is smaller than a predetermined size, storing said second packet in said hybrid storage area.</p>	<p>(5:27-28) "In the fast path case, network data is given to the host after the headers have been processed and stripped."</p> <p>(13:17-21) "As mentioned above, the fast-path flow puts a header into a header buffer that is then forwarded to the host. The host uses the header to determine what further data is following, allocates the necessary host buffers, and these are passed back to the INIC via a command to the INIC. The INIC then fills these buffers from data it was accumulating on the card..."</p> <p>(13:22-24) "This ... puts all the data into the header buffer or, if the header buffer is too small, uses a large (2K) host buffer for all the data."</p> <p>(12:44 to 13:4-5) "In order to receive 1514 byte frames (maximum ether frame size), however, we can only fit 2 buffers in a 4k page, which is not a substantial savings. Fortunately, network frames tend to be either large (~1500 bytes), or small (<256 bytes). We take advantage of this fact by allocating large and small receive buffers. If a received frame fits in a small buffer, the INIC will use a small buffer. Otherwise it will use a large buffer."</p> <p>(21: 20-22) "For segments less than a full 1460 byte payload, all of the received segment will be forwarded; it will be absorbed directly by the TDI client without any further MDL exchange."</p> <p>(7:28-38) "Clearly there are circumstances in which this does not make sense. When a small amount of data (500 bytes for example), with a push flag set indicating that the data must be delivered to the client immediately, it does not make sense to deliver some of the data directly while waiting for the list of addresses to DMA the rest. Under these circumstances, it makes more sense to deliver the 500 bytes directly to the host, and allow the host to copy it into its final destination. While various ranges are feasible, it is currently preferred that anything less than a segment's (1500 bytes) worth of data will be delivered directly to the host, while anything more will be delivered as a small piece (which may be 128 bytes), while waiting until receiving the destination memory address before moving the rest."</p>
<p>31. A method of transferring a packet from a communication link to a host computer, comprising:</p> <p>parsing a first packet received from a communication link to determine if the first packet conforms to a predetermined set of communication protocols;</p>	<p>(6:36-38) "The header is fully parsed by hardware and its type is summarized in a single status word. The checksum is also verified automatically in hardware, and a hash key is created out of the IP addresses and TCP ports to expedite CCB lookup."</p> <p>(6:30-32) "When a frame is received by the INIC, it must verify it completely before it even determines whether it belongs to one of its CCBs or not. This includes all header validation (is it IP, IPV4 or V6, is</p>

<p>generating, from identifiers of a source and a destination of the first packet extracted from the first packet, a flow key to identify a communication flow comprising the first packet;</p>	<p>the IP header checksum correct, is the TCP checksum correct, etc)."</p> <p>(6:10-14) "CCBs are initialized by the host during TCP connection setup."</p> <p>(6:10-12) "A CCB is a structure that contains the entire context associated with a connection. This includes the source and destination IP addresses and source and destination TCP ports that define the connection."</p> <p>(4:16-19) "The ... context is ... identified by the IP source and destination addresses and TCP source and destination ports."</p>
<p>updating a flow database to track a status of the communication flow;</p>	<p>(86:7-8) "If a match is found, then the frame will be processed against the CCB by the INIC."</p> <p>(6:10-16) "A CCB is a structure that contains the entire context associated with a connection. This includes the source and destination IP addresses and source and destination TCP ports that define the connection. It also contains information about the connection itself such as the current send and receive sequence numbers..."</p>
<p>associating a first operation code with the first packet to indicate whether the first packet is re-assembleable with another packet in the communication flow;</p>	<p>(85:26-30) "As frames are received, they are placed into 2K-byte DRAM buffers by the Receive hardware sequencer, along with 16 bytes of the above frame status. A pointer to the last byte + 1 of this buffer is queued into the Q_RECV queue. The pointer contains a bit (bit 29) that informs the microcode if this frame is definitely not a fast-path candidate..."</p> <p>(85:38) "If bit 29 is set, this frame is going slow-path."</p>
<p>maintaining a plurality of storage areas for transferring packets to a host computer, including:</p>	<p>(9:22-30) "We will instead write receive buffer addresses to the INIC as receive buffers are filled. In order to avoid having to write to the INIC for every receive frame, we instead allow the host to pass off a pages worth (4k) of buffers in a single write."</p> <p>2.3.2. Support small and large buffers on receive</p> <p>In order to reduce further the number of writes to the INIC, and to reduce the amount of memory being used by the host, we support two different buffer sizes. A small buffer contains roughly 200 bytes of data payload, as well as extra fields containing status about the received data bringing the total size to 256 bytes. We can therefore pass 16 of these small buffers at a time to the INIC. Large buffers are 2K in size."</p>
<p>a first storage area configured to store data portions of packets in the communication flow;</p>	<p>(9:31-33) "Note that when we have a large fast-path receive, a small buffer will be used to indicate a small piece of the data, while the remainder of the data will be DMA'd directly into memory."</p> <p>(10:25-32) "Meanwhile, the NetBIOS client will allocate a memory pool large enough to hold the entire NetBIOS message, and will pass this address or set of addresses down to the transport driver. The transport driver will allocate an INIC command buffer, fill it in with the list of</p>

<p>a second storage area configured to store a packet having an associated operation code indicating that the packet is not re-assemblable; and</p> <p>a third storage area configured to store packets less than a predetermined size and headers of packets having data portions stored in said first storage area; and</p> <p>storing the first packet in one or more of said plurality of storage areas for transfer to the host computer.</p>	<p>addresses, set the command type to tell the INIC that this is where to put the receive data, and then pass the command off to the INIC by writing to the command register. When the INIC receives the command buffer, it will DMA the remainder of the NetBIOS data, as it is received, into the memory address or addresses designated by the host."</p> <p>(11:2-6) "If the INIC receives a frame that does not contain a TCP segment for one of its CCB's, it simply passes it to the host as if it were a dumb NIC. If the frame fits into a small buffer (~200 bytes or less), then it simply fills in the small buffer with the data and notifies the host. Otherwise it places the data in a large buffer, writes the address of the large buffer into a small buffer, and again notifies the host."</p> <p>(13:4-5) "If a received frame fits in a small buffer, the INIC will use a small buffer. Otherwise it will use a large buffer."</p> <p>(13:17-18) "As mentioned above, the fast-path flow puts a header into a header buffer that is then forwarded to the host."</p> <p>See above.</p>
<p>32. The method of claim 31, wherein said storing comprises:</p> <p>selecting one or more storage areas of said plurality of storage areas on the basis of said operation code.</p>	<p>(85:26-30) "As frames are received, they are placed into 2K-byte DRAM buffers by the Receive hardware sequencer, along with 16 bytes of the above frame status. A pointer to the last byte + 1 of this buffer is queued into the Q_RECV queue. The pointer contains a bit (bit 29) that informs the microcode if this frame is definitely not a fast-path candidate..."</p> <p>(85:38) "If bit 29 is set, this frame is going slow-path."</p>
<p>33. The method of claim 31, wherein each of said plurality of storage areas is substantially equal in size to one memory page of the host computer.</p>	<p>(9:22-37) "We will instead write receive buffer addresses to the INIC as receive buffers are filled. In order to avoid having to write to the INIC for every receive frame, we instead allow the host to pass off a pages worth (4k) of buffers in a single write.</p> <p>2.3.2. Support small and large buffers on receive</p> <p>In order to reduce further the number of writes to the INIC, and to reduce the amount of memory being used by the host, we support two different buffer sizes. A small buffer contains roughly 200 bytes of data payload, as well as extra fields containing status about the received data bringing the total size to 256 bytes. We can therefore pass 16 of these small buffers at a time to the INIC. Large buffers are 2k in size...Since large buffers are 2k in size they are passed to the INIC 2 buffers at a</p>

	<i>time."</i>
<p>35. The network interface of claim 29, further comprising:</p> <p>a code generator configured to generate a code for the first packet to indicate whether the first packet is re-assemblable with a second packet in the second communication flow.</p>	<p>(85:26-30) "As frames are received, they are placed into 2K-byte DRAM buffers by the Receive hardware sequencer, along with 16 bytes of the above frame status. A pointer to the last byte + 1 of this buffer is queued into the Q_RECV queue. The pointer contains a bit (bit 29) that informs the microcode if this frame is definitely not a fast-path candidate..."</p> <p>(85:38) "If bit 29 is set, this frame is going slow-path."</p>
<p>36. The network interface of claim 29, further comprising:</p> <p>a code generator configured to generate a code for the first packet to indicate whether the first packet is larger than the predetermined size.</p>	<p>(162:6-10) "27:00 address Represents the last address +1 to which frame data was transferred. This can be used to determine the size of the frame received by first subtracting one from the address then comparing with the buffer boundary as indicated by the size bits."</p> <p>(165:20) "14:00 LengthCnt Refer to E110 Technical Manual."</p>
<p>37. A communication interface configured to store packets for transfer to a host computer, comprising:</p> <p>a parser configured to determine whether a packet includes a data portion;</p> <p>a re-assembly storage area configured to store data portions of a plurality of packets from a single communication flow; and</p>	<p>(6:36-38) "The header is fully parsed by hardware and its type is summarized in a single status word. The checksum is also verified automatically in hardware, and a hash key is created out of the IP addresses and TCP ports to expedite CCB lookup."</p> <p>(6:30-32) "When a frame is received by the INIC, it must verify it completely before it even determines whether it belongs to one of its CCBs or not. This includes all header validation (is it IP, IPV4 or V6, is the IP header checksum correct, is the TCP checksum correct, etc)."</p> <p>(9:31-33) "Note that when we have a large fast-path receive, a small buffer will be used to indicate a small piece of the data, while the remainder of the data will be DMA'd directly into memory."</p> <p>(10:25-32) "Meanwhile, the NetBIOS client will allocate a memory pool large enough to hold the entire NetBIOS message, and will pass this address or set of addresses down to the transport driver. The transport driver will allocate an INIC command buffer, fill it in with the list of addresses, set the command type to tell the INIC that this is where to put the receive data, and then pass the command off to the INIC by writing to the command register. When the INIC receives the command buffer, it will DMA the remainder of the NetBIOS data, as it is received, into the memory address or addresses designated by the host."</p>

<p>a hybrid storage area configured to store:</p> <p>header portions of the plurality of packets; and</p> <p>one or more packets smaller than a first predetermined size.</p>	<p>(13:17-18) "As mentioned above, the fast-path flow puts a header into a header buffer that is then forwarded to the host."</p> <p>(13:4-5) "If a received frame fits in a small buffer, the INIC will use a small buffer. Otherwise it will use a large buffer."</p>
<p>39. The communication interface of claim 37, further comprising:</p> <p>a code generator configured to generate a code for each packet received at the communication interface;</p> <p>wherein said code is configured to indicate a type of storage area in which said packet may be stored.</p>	<p>(85:26-30) "As frames are received, they are placed into 2K-byte DRAM buffers by the Receive hardware sequencer, along with 16 bytes of the above frame status. A pointer to the last byte + 1 of this buffer is queued into the Q_RECV queue. The pointer contains a bit (bit 29) that informs the microcode if this frame is definitely not a fast-path candidate..."</p> <p>(85:38) "If bit 29 is set, this frame is going slow-path."</p>
<p>40. The communication interface of claim 37, wherein said parser is further configured to determine a size of each packet received at the communication interface.</p>	<p>(162:6-10) "27:00 address Represents the last address + 1 to which frame data was transferred. This can be used to determine the size of the frame received by first subtracting one from the address then comparing with the buffer boundary as indicated by the size bits."</p> <p>(165:20) "14:00 LengthCnt Refer to E110 Technical Manual."</p>

As demonstrated above, the '296 app. provides full support for pending claims 1-40. The disclosure of the '296 app. has been included in each of the applications in the continuous chain up to and including the present application. Thus, claims 1-40 are entitled to the benefit of the August 27, 1998, filing date of the '296 app.

The fact that applicants did not use the exact same nonce words in their disclosure as were used by Muller to claim the subject matter that applicants invented is no reason under the law to deny applicants the benefit of the filing date of the '296 app. Stated differently, while Muller may have coined the nonce words "hybrid storage area" and "re-assembly storage area," Muller did not invent the claims that recite those terms. Indeed, while applicants copied the claims of Muller to provoke an interference, Muller copied applicants' inventions and obtained a patent on them. Muller should not be

allowed to succeed with this theft simply by creating nonce words that applicants did not use in their earlier disclosure of the same subject matter.

II. 35 U.S.C. 102

A. Muller

Claims 1-40 stand rejected under 35 U.S.C. 102(e) as being clearly anticipated by Muller.

As discussed above, applicants respectfully assert that Muller is not prior art to any of those claims.

B. Thompson

Claims 17, 34 and 38 stand rejected under 35 U.S.C. 102(b) as being anticipated by U.S. Patent No. 5,917,828 to Thompson. Regarding claim 17, the Office Action states:

Regarding claim 17, Thompson disclosed a method of network communication, the method comprising: providing a computer having a processor and a memory, the memory including first, second and third storage areas that are accessible by a communication interface (Fig. 11, memories 1108, 1122, and 1140);

receiving, by the communication interface, a first packet, and storing the first packet in the first storage area (Thompson, Fig. 11, All packets received are first stored in cell FIFO 1108);

receiving, by the communication interface, a second packet, storing a header of the second packet in the first storage area (Thompson, Fig. 11, All packets received are stored in cell FIFO 1108), and storing data of the second packet in the second storage area (Thompson, col. 8, lines 27-31, Thompson disclosed data from the received packets being stored in local memory buffer 1122 in certain conditions); and

receiving, by the communication interface, a third packet, and storing data of the third packet in the third storage area, the third storage area containing data from a plurality of packets and corresponding to an application running on the computer, the third storage area containing no headers (Thompson, col. 8, lines 40-45, Thompson disclosed the contents of the local memory buffer, which include the payloads of received packets, to be stored in main memory buffer 1140).

Applicants respectfully disagree with this rejection. Applicants respectfully assert that it is not at all clear that Thompson discloses the recitation of "the third storage area

containing data from a plurality of packets and corresponding to an application running on the computer, the third storage area containing no headers.” Indeed, the Office Action explanation that: “Thompson disclosed the contents of the local memory buffer, which include the payloads of received packets, to be stored in main memory buffer 1140,” does not show that Thompson teaches “the third storage area containing no headers,” as recited in claim 17. That is, although the Office Action states that “the contents of the local memory buffer... *include* the payloads of received packets,” that is not the same as saying that there are “no headers” in those contents.

Moreover, a closer analysis of Thompson reveals that the main memory buffer does include packet headers. For example, in column 7, lines 47-57, Thompson states:

If the PDU is slightly larger than one memory buffer (4096+492 bytes), the status will be combined in local buffer memory 1006 with the pointer to the local memory buffer plus the remaining data bytes and this will be written as one burst to the status queues in main memory 1008. This is important because most large PDU transfers are equal to the memory page size plus a small TCP/IP or UDP/IP header and would normally use only a small amount of a second host memory buffer. This feature also combines the status write and the write of the overflow data into one larger burst write from local memory 1006 to host memory 1008.

In other words, “a small TCP/IP or UDP/IP header” is transferred from the local memory buffer to the main memory buffer according to Thompson. While this header may fit into one or two of the ATM payloads, it is certainly a header. Moreover, there is no indication in the Office Action of how one of ordinary skill in the art would have modified Thompson to meet the limitation recited in claim 17 of “the third storage area containing no headers.”

For at least these reasons, applicants respectfully assert that Thompson does not anticipate or render obvious claim 17 or any claim that depends from claim 17, such as claims 34 or 38.

III. Conclusion

Applicants have responded to each item of the Office Action to demonstrate that the pending claims are in condition for allowance. In particular, applicants have shown that the pending claims have support in the ’296 app. and a continuous chain of

applications up to and including the present, so that Muller is not prior art to any of the claims. Should the Examiner have any question about this response or this application, he is respectfully requested to telephone the undersigned.

Respectfully submitted,

/Mark Lauer/
Mark Lauer
Reg. No. 36,578
6601 Koll Center Parkway
Suite 245
Pleasanton, CA 94566
Tel: (925) 621-2121
Fax: (925) 621-2125